arXiv:1412.2446v2 [math.GT] 10 May 2016

AN UPPER BOUND FOR THE VOLUMES OF COMPLEMENTS OF PERIODIC GEODESICS

MAXIME BERGERON, TALI PINSKY, AND LIOR SILBERMAN

Abstract. A periodic geodesic on a surface has a natural lift to the unit tangent bundle; when the complement of this lift is hyperbolic, its volume typically grows as the geodesic gets longer. We give an upper bound for this volume which is linear in the geometric length of the geodesic.

1. Introduction

A closed curve on a surface S can be naturally lifted to the unit tangent bundle T^1S by traversing the curve in a chosen direction and associating to each point on the curve its tangent direction (note that there are two lifts, corresponding to the two directions along the curve). Such a lift is an embedding of $\mathbb{S}^1$ into the 3-manifold T^1S and, considering its ambient isotopy class, we obtain a knot.

If we equip S with a hyperbolic Riemannian metric, the isotopy class of each (non-peripheral and non-trivial) closed curve contains a periodic geodesic (the representative of shortest length). The associated knot is then a closed orbit of the geodesic flow on T^1S which does not depend on the chosen metric. Such periodic geodesics have long been an object of interest from the topological [10, 18] and dynamical [2] points of view. However, except for a few promising results for closed geodesics on the modular surface [8, 11, 13, 16] and other triangle groups [8], they have essentially never been studied as knots. One of the reasons for this is that these knots are embedded in T^1S which, unless S happens to be the modular surface, cannot be embedded into $\mathbb{S}^3$. As such, there is a limited supply of knot theoretic invariants available to try and characterize them.

Nevertheless, the complement of a lift of a periodic geodesic in the unit tangent bundle is a hyperbolic 3-manifold as soon as the curve is filling [9] in the surface. In this case, its volume is a topological invariant and we give the first known estimate for it:

Date: May 11, 2016.

M.B. was supported by an NSERC Alexander Graham Bell CGS-D Scholarship.

L.S. was partly supported by an NSERC Discovery Grant.

Theorem 1.1. *Let $S = \mathbb{H}^2/\Gamma$ be a hyperbolic surface. There is a constant $C_S > 0$ such that, for any finite set γ of (primitive) periodic geodesics on S, we have*

$$\mathrm{Vol}(T^1 S \setminus \gamma) \leq C_S \, |\gamma| \, . \tag{1}$$

Here, $|\gamma|$ denotes the total length of the geodesics in γ.

Remark. The length $|\gamma|$ is measured with respect to the fixed hyperbolic metric on S. We emphasize, however, that the volume is topological as is the linearity of the bound. Changing the metric will only affect the constant.

Our proof of Theorem 1.1 begins with the modular surface M. In this case, there is a symbolic description of the closed geodesics: they can be coded by their continued fraction expansion as done by E. Artin [3, 17]. On the other hand, this symbolic description was recently shown by Ghys [11] to have a topological interpretation. The key to our linear bound for M lies in the connection between a decomposition of $T^1 M$ and a complexity measure for the codes of the geodesics in γ obtained via Ghys's explicit *template* (c.f. Section 2).

Remark. In Example 3.1 we exhibit a family of arbitrarily long geodesics on the modular surface whose associated knot complements have uniformly bounded volume. Nevertheless, numerical evidence [5] indicates that, in some situations, the growth is linear in the geometric length of the geodesics.

Remark. It is interesting to compare our result for the modular surface with Guéritaud and Futer's volume estimates for once-punctured torus bundles and two-bridge link complements [12]. They use the same symbolic coding for elements in $SL_2(\mathbb{Z})$. However, their volume bounds are linear in the period of the continued fraction expansion, whereas ours are linear in the geometric length (c.f. Bridgeman [6]).

The paper unfolds as follows: In Section 2, we review the coding of geodesics on the modular surface M by positive words. In Section 3, we obtain a volume bound in terms of the combinatorics of the coding, from which we deduce Theorem 1.1 for M. Finally, in Section 4, we deduce the general case of Theorem 1.1 by relating covers of punctured surfaces to covers of M.

We are grateful to Juan Souto for suggesting that we study this invariant and for offering useful insights. We would also like to thank Yair Minsky, Jessica Purcell and Kasra Rafi for helpful discussions.

2. Background: Coding of geodesics on the modular surface

In this section, we review Ghys's construction [11] of a template and the associated coding for the geodesic flow on the modular surface.

We start with the isomorphism $\mathrm{PSL}_2(\mathbb{Z}) \simeq C_2 * C_3$. For later reference, we fix elements κ_0 and $\omega \in \mathrm{PSL}_2(\mathbb{Z})$ of order 2 and 3, respectively, so that $\mathrm{PSL}_2(\mathbb{Z}) = \langle \kappa_0 \rangle * \langle \omega \rangle$. Restricting every homomorphism to the generating set, we see that $\mathrm{PSL}_2(\mathbb{Z})$ has an essentially unique surjection onto C_3, and hence a unique normal subgroup Γ_3 of index 3. We write $M_3 := \mathbb{H}^2/\Gamma_3$ for the associated hyperbolic manifold, the unique normal three-fold cover of M, depicted in Figure 1.

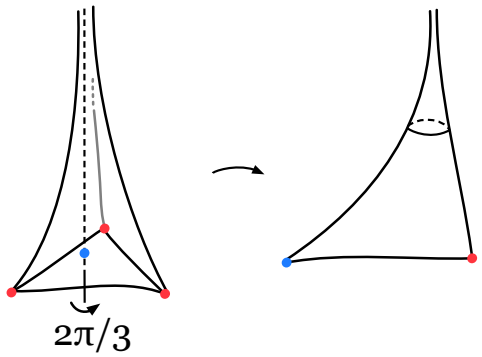

Figure 1. The three-fold cover M_3 of the modular surface M.

Remark. For an alternative construction of Γ_3, note that $\mathrm{PSL}_2(\mathbb{F}_3) \simeq S_4$ (consider the action on $\mathbb{P}^1(\mathbb{F}_3)$) and that the 2-Sylow subgroup of S_4 is normal of index 3; this gives a surjection onto C_3 and shows that the kernel contains $\Gamma(3)$ so that Γ_3 is a congruence subgroup.

Let $p \in \mathbb{H}^2$ be the point fixed by ω and notice that, since $\mathrm{Stab}_{\mathrm{PSL}_2(\mathbb{Z})}(p)$ coincides with $\langle \omega \rangle$, p is not a fixed point of an elliptic element of Γ_3. It follows that the associated Dirichlet domain U_3 (the set of points of $\mathbb{H}^2$ closer to p than to any other point of the orbit $\Gamma_3 \cdot p$) is a fundamental domain for M_3. In fact, as shown in Figure 2, U_3 is an ideal triangle. Moreover, κ_0 fixes a point q_0 along one of the sides of this triangle which we label J_0, and its conjugates $\kappa_i = \omega^i \kappa_0 \omega^{-i}$ fix points $q_i = \omega^i q_0$ along the other two sides J_1 and J_2. For convenience (as shown in the figure) we choose our identification of $\mathbb{H}^2$ with the disc so that p is the centre of the disc. In that case ω acts by rotation by the angle $2\pi/3$,

cyclically permuting the vertices of the triangle, the arcs J_i connecting them, and the elliptic fixed points q_0, q_1 and q_2.

We now use this picture to study certain geodesics on M. Observe first that any set γ of periodic geodesics lifts to a set $\tilde{\gamma}$ of closed curves in M_3 (each periodic geodesic in M lifts either to three periodic geodesics of the same length, or to a single geodesic of three times the length, but this is immaterial for our arguments). Moreover, any geodesic in M_3 has a lift to an (infinite) geodesic on $\mathbb{H}^2$ connecting two points on the boundary. This lift may be chosen to cross any particular fundamental domain for $\mathrm{PSL}_2(\mathbb{Z})$, specifically U_3. If the geodesic is periodic, its ends cannot lie on the cusps of $\mathrm{PSL}_2(\mathbb{Z})$ and, in particular, on the vertices of U_3. Hence, acting by the element of order 3, we may choose the lift to start at $I_0 \subset \partial\mathbb{H}^2$, so that it will enter the triangle U_3 through its side J_0.

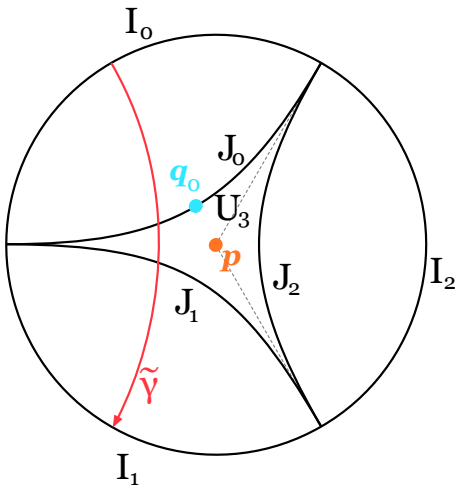

FIGURE 2. A fundamental domain, U_3, for a three-fold cover M_3 of the modular surface.

We have seen that it is enough to encode geodesics through U_3 starting at I_0. Accordingly, let $\tilde{\gamma}$ be such an infinite geodesic. We now construct a different (disconnected) lift of γ, consisting of a sequence of segments in U_3. The first part of this new lift will be the segment of $\tilde{\gamma}$ between its entry point to U_3 along J_0 and its exit point along either J_1 or J_2. We begin our code by x or y depending on the two possibilities.

Suppose our segment ends on J_i. Acting by κ_i, the remaining ray of $\tilde{\gamma}$ (the part after the end of the segment) now *begins* on J_i (but usually

not on the point where the segment ended). Rotating by a power of ω, we may assume instead that the remaining ray enters again via J_0. It will then exit via one of the other sides and once again: we break off a segment, record a letter x or y, apply κ_1 or κ_2 and a rotation, and continue as before.

If $\tilde{\gamma}$ projects to a periodic geodesic γ on M, the resulting infinite word will be periodic. In that case the *code* w_γ for γ will be the the primitive part of this periodic word. We write n_γ for the number of (cyclic) subwords of the form xy in w_γ. This corresponds to the period in the continued fraction expansion of γ. In more generality, for $\tilde{\gamma} = \cup_{i=1}^{k} \tilde{\gamma}_i$ (a set of geodesics projecting to a collection of periodic geodesics γ on M) we define n_γ to be the sum $n_\gamma = \sum_{i=1}^{k} n_{\gamma_i}$. It is shown in [11] that, up to cyclic permutation, the word w_γ only depends on the projection of $\tilde{\gamma}$ to M and, conversely, that any such word encodes a periodic geodesic.

As noted in the introduction, this coding also arises from the continued fraction expansion of a primitive modular geodesic (c.f. Series [17]). There, n_γ is exactly half the period of the (even) continued fraction corresponding to γ. Constructing the code as above, however, yields a natural interpretation in terms of a *template* embedded in T^1M whose existence is due to Birman and Williams [4] and which was explicitly constructed by Ghys [11]. This template is a branched surface with boundary, equipped with a semi-flow such that any finite set $\tilde{\gamma}$ of geodesics may be deformed into closed flow orbits on the template by an ambient isotopy (of $\tilde{\gamma}$ together with the tangent direction). Ghys [11] showed that the template can be chosen to have a single branch-line lying over J_0 in T^1M together with two bands such that, after the isotopy, one contains the geodesic segments passing in U_3 from J_0 to J_1 while the other contains the segments passing from J_0 to J_2. From this point of view the coding given above describes the sequence by which the image of $\tilde{\gamma}$ under the isotopy travels through the two template bands. Figure 3 depicts the template with its embedding in T^1M, i.e., the complement of a trefoil knot in $\mathbb{S}^3$.

Since our first upper bound on the volume of the complement of the lift of γ in T^1M will be obtained in terms of the combinatorics of w_γ, we will need the following relation between this codeword and the length of γ:

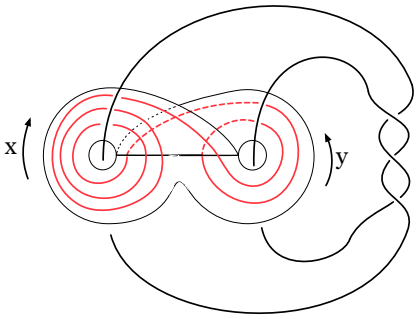

FIGURE 3. Ghys's template for the modular surface embedded in its unit tangent bundle (the complement of a trefoil in $\mathbb{S}^3$) along with a periodic orbit corresponding to the code word x^3y^2.

Lemma 2.1. *The length L of the closed geodesic γ coded by the word $x^{k_1}y^{m_1}\dots x^{k_{n_\gamma}}y^{m_{n_\gamma}}$ (recall that $k_i, m_i \in \mathbb{Z}_{\geq 1}$) satisfies*

$$L \geq \sum_{i=1}^{n_\gamma} \left(\log(k_i) + \log(m_i)\right), \text{ and}$$
$$L \geq (\log 2) n_\gamma.$$

Proof. Letting $x = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $y = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, the geodesic γ corresponds to the conjugacy class in $\mathrm{PSL}_2(\mathbb{Z})$ of the element $\prod_{i=1}^{n_\gamma} x^{k_i}y^{m_i}$. We note that

$$x^k y^m = \begin{pmatrix} km+1 & k \\ m & 1 \end{pmatrix}.$$

Since the entries of a product of matrices are monotone in the entries of the factors (when the entries are non-negative) we see that

$$\gamma = \prod_i x^{k_i}y^{m_i} \geq \begin{pmatrix} \prod_i (k_i m_i + 1) & 0 \\ 0 & 1 \end{pmatrix}.$$

The entries of this matrix are connected to the length by $\mathrm{Tr}(\gamma) = e^L + e^{-L}$, so that $e^L \geq \mathrm{Tr}(\gamma) - 1 \geq \prod_i (k_i m_i + 1)$. The first claim follows immediately, and for the second we note that $k_i m_i + 1 \geq 2$. □

3. Triangulation of link complements and the combinatorial upper bound

Let γ be a (filling) collection of periodic geodesics on M, let w_γ be their coding and consider their lifts $\tilde{\gamma}$ in T^1M, isotoped onto Ghys's

template as in the previous section. Proceeding to drill out $\tilde{\gamma}$ from T^1M, we obtain a hyperbolic 3-manifold. We will estimate the volume of this manifold by constructing a triangulation. However, the number of tetrahedra in our triangulation is proportional to the word length of w_γ, so directly using the boundedness of volumes of hyperbolic tetrahera does not give a useful bound. The following example illustrates the problem:

Example 3.1. Let γ be a geodesic on M coded by the word $x^n y^m$ for some natural numbers n and m. Its length is roughly proportional to $\log(n+m)$ and, in particular, tends to infinity with n and m. The number of crossings on the template (which will be proportional to the number in tetrahedra in our triangulation) is then $n+m$. On the other hand, these geodesics correspond to a knot winding more and more around the trefoil, first in one template ear and then in the other. Thus, their volumes are all bounded by the volume of $T^1M \setminus (\gamma \cup \alpha \cup \beta)$ where γ is the geodesic corresponding to xy, and α and β are both trivial knots encircling one strand of γ and the strand of the trefoil in the centre of the corresponding ear (c.f Adams [1]). Note that $n_\gamma = 1$ for any geodesic γ in this family.

In order to circumvent this difficulty, in Section 3.2 we will establish that large groups of tetrahedra in our triangulation share an edge. The volume of a tetrahedron can be expressed in terms of its dihedral angles via the so-called Lobachevsky function, and we will analyze it to show that the sum of the volumes of the tetrahedra in the group grows sublinearly in the number of tetrahedra. This will allow us to attain our linear bound for M.

Remark. With the hope that no confusion will arise, we will no longer distinguish a collection γ of geodesics on M from its lift $\tilde{\gamma}$ to T^1M. Both of them will be henceforth denoted by γ.

3.1. **Triangulation.** Let us fix once and for all a finite collection γ of periodic geodesics on M. In order to triangulate $T^1M \setminus \gamma$, we first decompose it into pieces of three types (A, B and C) along with a "remainder" (D). There will be two pieces of type A, two pieces of type B, four pieces of type C and one piece of type D. As such, writing $\mathrm{Vol}(A)$ (respectively $\mathrm{Vol}(B)$, $\mathrm{Vol}(C)$ and $\mathrm{Vol}(D)$) for an upper bound on the volume of a piece of type A (respectively B, C and D), the volume of T^1M will then be bounded from above by

$$2\,\mathrm{Vol}(A) + 2\,\mathrm{Vol}(B) + 4\,\mathrm{Vol}(C) + \mathrm{Vol}(D). \tag{2}$$

3.1.1. *Cutting.* Considering Figure 3 once again, we define some cutting segments on the template that will be used to decompose $T^1M \setminus \gamma$. First, follow the flow-lines up from the critical point into both template ears. After propagating up past the branch-line, stop at some point p in the x ear and some point q in the y ear. Then, connect each of these points to the outer boundary of the template ear by two additional segments as shown in Figure 4.

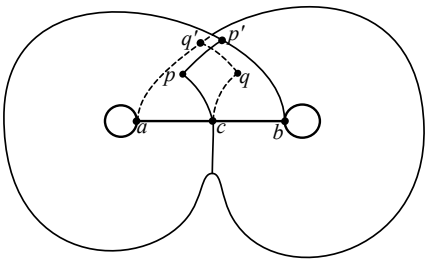

FIGURE 4. The cutting segments on the template; the segments connecting p and q to the critical point c are chosen to lie along the flow-lines.

3.1.2. *Pieces.* We can now describe the four different types of pieces shown in Figure 5. All relative position adjectives refer to the figure.

Type A**:** Choose a point o_1 on the trefoil, above the centre of the x ear, and use it to cone-off the (upper side of the) x ear without including the quadrilateral $cbp'p$. This is the first piece of type A, we call it A_1. The second piece of type A is the cone of the bottom side of the y ear with a point o_1' on the trefoil below the centre of the y ear. We call it A_2.

Type B**:** Consider the prism obtained by taking the product of the y ear without the quadrilateral $qq'ac$ with the segment cp as in Figure 5(b). The quadrilateral $cbp'p$ is contained in this cylinder as $cb \times cp$ and the upper face of this cylinder is bounded by an edge connecting p to itself along the outer side and an edge connecting p' to itself along the inner side. Choose a point o_2 on the trefoil above the centre of the y ear and cone-off the upper face of this cylinder except for the segment pp' which is coned-off to o_2 by two $pp'o_2$ triangles: one on each sides of the $\tilde{\gamma}$ strands coming into the y-ear. This "coned-off" cylinder is the first piece of type B, we call it B_1. The second piece of this

type is on the bottom side of the x ear, connected to a point o_2' on the trefoil below the centre of the x ear. We call it B_2.

Type C: The four pieces of type C are tetrahedra used to close the gaps between pieces of type A and type B: $pp'o_1o_2$ connects B_1 to A_1, $pp'o_2o_2'$ connects B_1 to B_2, $qq'o_1'o_2'$ connects B_2 to A_2 and $qq'o_2o_2'$ connects B_2 to B_1.

Type D: There is only one piece of type D, it consists of the remainder of the complement of the trefoil in $\mathbb{S}^3$.

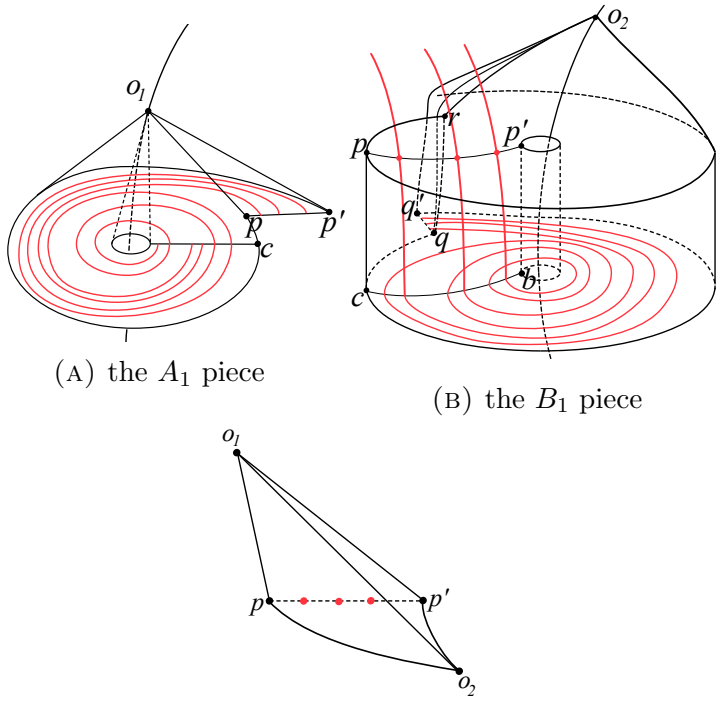

(A) the A_1 piece

(B) the B_1 piece

(C) the C_1 piece

FIGURE 5

These pieces fit together as follows (as usual, relative position adjectives refer to Figure 5):

(i) Piece A_1 is glued to piece B_2 along the x ear. It is glued to a C piece that we will call C_1 along the triangle $pp'o_1$. The remainder of its coning face is glued to the D piece.

(ii) Piece B_1 is glued to piece A_2 along the y ear. Recall that B_1 has two $pp'o_2$ faces resulting from the double coning along pp'. It is glued to C_1 (connecting it to A_1) along its front $pp'o_2$ face and it is glued to another C piece which we call C_2 (connecting it to B_2)

along its back $pp'o_2$ face. It is glued to piece B_2 along the square $pcqr$ and it is glued to yet another C piece which we call C_3 along $qq'o_2$ (connecting it to B_2). The remaining outer faces of piece B_1 are glued to the D piece.

(iii) The situation for piece A_2 (respectively piece B_2) is analogous to that of piece A_1 (respectively piece B_1).

(iv) Each C piece is connected to either A_i and B_i or B_i and B_j. It is glued to the D piece along its two faces that do not intersect γ.

3.1.3. *Pre-triangulation.* Our next step is to triangulate the exterior faces of the pieces. This "pre-triangulation" will then be completed to a triangulation of the interiors. Gluing the pieces along triangles will ensure the compatibility of our volume estimates.

We start with piece A_1. There, we deem the points o_1, p, p', a, c as well as all intersections of pp' and ac with γ to be vertices. After adding the edges indicated in Figure 5, our next set of edges consists of all resulting subsegments of ac and pp'. The vertices along pp' corresponding to intersections with γ are then all joined by edges to o_1. This turns the face $pp'o_1$ on which A_1 is glued to C_1 into a union of triangles. Moreover, the faces contained in the x ear of the template band along which piece A_1 is glued to piece B_2 turn into either bigons, triangles or higher valence polygons as shown in Figure 6. In the latter cases, we fix any triangulation of the polygon within the template ear. In the case of a bigon, the two parallel edges between two γ segments on a face of A_1 can be collapsed along the γ segments in the direction of the template semi-flow, resulting in a single edge along pp'. The exterior faces of piece A_2 are triangulated in an analogous manner.

To treat piece B_1, observe that its face contained in the y ear of the template is already triangulated by the choices and identifications made for the corresponding face of piece A_2. After adding the edges indicated in Figure 5, to triangulate its remaining faces, we first add edges connecting all intersection points of γ with qq' to o_2 and then add edges connecting all intersection points of γ with pp' to o_2 in two different ways – one in each of the two $pp'o_2$ faces on each side of the strands coming in from the x-band of the template. Finally, we add a diagonal edge pq to the square face $prqc$. The exterior faces of piece B_2 are triangulated in an analogous manner.

To triangulate the exterior faces of pieces of type C, we simply connect each ideal vertex to the points o_1 and o_2 or o_1' and o_2'.

3.2. **A bound for the modular surface.** Recall that there is a uniform upper bound for the volume of geodesic hyperbolic tetrahedra. In particular, we can bound the volume of any piece by the number

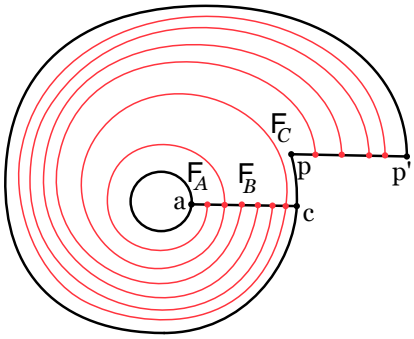

FIGURE 6. Different polygons induced on the x ear of the template band. Here, F_A is a triangle, F_B is a pentagon that would be subdivided into 3 triangles, F_C is a square and the other faces are all bigons which would collapsed into single edges along pp'.

of tetrahedra it contains. In this section, we use this observation to establish Theorem 1.1 for the modular surface in two steps. First, we bound the number of tetrahedra in pieces of type A, C and D by n_γ. This is sufficient for our purposes since, by Lemma 2.1, upper bounds linear in n_γ are also linear in $|\gamma|$. Unfortunately, our triangulation of pieces of type B will require a number of tetrahedra that is quadratic in n_γ. Accordingly, our second step uses the Lobachevsky function to show that we still obtain a linear bound in $|\gamma|$.

Proposition 3.2. *The volume of a piece of type A, C or D is at most linear in n_γ and, thus, is at most linear in $|\gamma|$.*

Proof. To begin, we complete our pre-triangulations of these pieces to triangulations of their interior. Since the piece of type D is a fixed polyhedron which doesn't intersect γ, we can endow it with an arbitrary fixed triangulation. To complete the pre-triangulation of pieces of type C, we add two edges joining each intersection of γ along pp' (respectively qq') to the other two vertices. This allows us to triangulate them using $n_\gamma + 1$ tetrahedra. The volume contributed by the C and D pieces is therefore at most linear in $|\gamma|$.

Unlike these first two cases, pieces of type A require a more careful count. Recall that a bigon in our pre-triangulation corresponded to a rectangle in the x (respectively y) ear of the template, bounded by two adjacent γ arcs and two "parallel" edges (see Figure 6). These bigons were eliminated by collapsing the parallel edges along the γ arcs in the direction of the template's semi-flow, resulting in a single edge

along pp' (respectively qq'). If this edge happened to bound a further bigon, it would once again be collapsed along adjacent γ arcs in the direction of the semi-flow. This process continued until the two γ arcs emanating from the collapsed edge ceased to be adjacent, i.e., their next intersection with the branch line occurred at two points separated by further intersection with γ coming from arcs in the y (respectively x) ear (see Figure 3). If n arcs coming from the y (respectively x) ear fall in between the endpoints of two adjacent γ arcs, we have an $n+2$-gon in the template which was triangulated using n triangles. We can then cone them off to o_1 (respectively o_1') using n tetrahedra. Since there are a total of n_γ arcs coming in from the y (respectively x) ear, the number of tetrahedra used to complete our triangulation of a piece of type A is at most linear in n_γ. This completes the proof of the Proposition. □

For our next gain, let

$$\Lambda(\theta) = -\int_0^\theta \log|2\sin u|\,\mathrm{d}u.$$

be the Lobachevsky function.

Lemma 3.3. *If T is a tetrahedron with dihedral angle α along the edge e, then*

$$\mathrm{Vol}(T) \leq 2\Lambda\left(\frac{\alpha}{2}\right).$$

Proof. We may assume T is an ideal tetrahedron, and let β, γ be the dihedral angles at the two other edges incident to a vertex at the end of e. As recounted in [15, Lemma 2], it was shown by Lobachevsky that $\alpha+\beta+\gamma=\pi$ and that

$$\mathrm{Vol}(T) = \Lambda(\alpha)+\Lambda(\beta)+\Lambda(\gamma).$$

It is verified in [12, Prop. 6.6] that the expression on the right is a concave function on the simplex $\alpha+\beta+\gamma=\pi$. In particular,

$$\mathrm{Vol}(T) \leq \Lambda(\alpha) + 2\Lambda\left(\frac{\pi-\alpha}{2}\right).$$

By the identity (c.f. [15, Lemma 1])

$$\Lambda(2\theta) = 2\Lambda(\theta) + 2\Lambda\left(\frac{\pi}{2}+\theta\right) = 2\Lambda(\theta) - 2\Lambda\left(\frac{\pi}{2}-\theta\right),$$

we see that $2\Lambda\left(\frac{\pi}{2}-\frac{\alpha}{2}\right) = 2\Lambda\left(\frac{\alpha}{2}\right) - \Lambda(\alpha)$ and the claim follows. □

Proposition 3.4. *Let $\{T_i\}_{i=1}^s$ be a family of $s \geq 2$ tetrahedra in $\mathbb{H}^3$, all sharing the same edge e. Then*

$$\sum_i \mathrm{Vol}(T_i) \leq C \log(s)$$

where C is a universal constant.

Proof. For $s \leq 5$ the claim follows from the uniform bound on the volumes of tetrahedra, so from now on we assume $s \geq 6$. Let α_i be the dihedral angle of the tetrahedron T_i at the edge e, so that the total angle $\sum_i \alpha_i \leq 2\pi$. By the Lemma, we have

$$\sum_i \mathrm{Vol}(T_i) \leq 2 \sum_i \Lambda\left(\frac{\alpha_i}{2}\right).$$

Since $\Lambda''(\theta) = -\cot(\theta)$ is non-positive on $[0, \pi/2]$ the Lobachevsky function is concave there and it follows that

$$\sum_i \mathrm{Vol}(T_i) \leq 2s\Lambda\left(\frac{\sum_i \alpha_i}{2s}\right) \leq 2s\Lambda\left(\frac{\pi}{s}\right),$$

where the second inequality follows from the monotonicity of Λ in $[0, \pi/6]$ and the fact that $\sum_i \alpha_i/2s \leq \pi/6$ (here we used the assumption $s \geq 6$).

Finally, by the absolutely convergent series expansion

$$\Lambda(\theta) = \theta\left(1 - \log|2\theta| + \sum_{k=1}^{\infty} \frac{B_k}{2k} \frac{(2\theta)^{2k}}{(2k+1)!}\right)$$

(c.f. [15, p. 18]) we have

$$x\Lambda(\frac{\pi}{x}) \leq c + \pi \log(x) \leq C \log(s)$$

for some constants $c, C > 0$. □

We are now ready to prove:

Theorem 3.5 (Theorem 1.1 for the modular surface)**.** *There is a constant $C > 0$ such that, for any finite collection γ of periodic geodesics on M, we have*

$$\mathrm{Vol}(T^1M \setminus \gamma) \leq C\ |\gamma|.$$

Proof. In light of Equation 2 and Proposition 3.2 it remains to consider the contribution of the pieces of type B and, by symmetry, it suffices to consider piece B_1.

Before completing our pre-triangulation of this piece into a triangulation, we introduce some notation and describe additional vertices within the piece's interior. Choose a line l parallel to the segment pp'

and slightly below it within the $pp'bc$ rectangle. There are n_γ intersections of γ with l, they correspond to the strands arriving from the x ear into the y ear. We label them from right to left as $p_1^l, p_2^l, \ldots, p_{n_\gamma}^l$ and add a vertex x_i between p_i^l and p_{i+1}^l. Next, we denote the intersection point of the strand starting at p_i^l with the branch-line by p_i and note that these two points correspond to the same point at infinity in our 3-manifold $T^1M \setminus \gamma$. Lastly, we denote by $[p_i, p_{i+1}]$ (respectively (p_i, p_{i+1})) the closed (respectively open) subinterval of the branch-line between p_i and p_{i+1}.

We now describe the edges of our triangulation. First, add any arc in the complement of the x_i's and p_i^l's within l to our set of edges. Second, whenever there is a γ segment starting in $[p_i, p_{i+1}]$ and reaching the interval (p_j, p_{j+1}) after winding around the y ear, we add an edge connecting x_i to x_j, running roughly above this γ segment parallel to the template and within the B piece. If there is a segment of γ emanating from $[p_i, p_{i+1}]$ and passing through qq' before returning to the branch-line, we connect x_i to o_2 by a "diagonal" edge above the previously added edges within the B piece. Note that some of these edges may correspond to many γ segments in the y ear of the template (the bottom face of the B piece) and that some of these edges may connect a point x_i to itself. Finally, we join x_i to all γ intersections within $[p_i, p_{i+1}]$ by an edge and connect all γ intersections in $[c, p_{n_\gamma}]$ to o_2 by an edge. We are now ready to formalize the remainder of our triangulation.

Consider adjacent intersection points of γ with an interval $[p_i, p_{i+1}]$. If the corresponding strands of γ split into different segments $[p_j, p_{j+1}]$ and $[p_k, p_{k+1}]$ after winding around the y ear, there is a $(k-j)$-gon on the template band which we have already triangulated (see Section 3.1.3). We also have a $2(k-j)+1$-gon in the plane cutting through the x_i's above the template band and vertical edges connecting the two polygons as illustrated in Figure 7. This polyhedron is easily triangulated using a number of tetrahedra that is linear in $k-j$. Since there are n_γ intervals $[p_i, p_{i+1}]$, the number of tetrahedra resulting from all such splittings is linear in n_γ.

On the other hand, if the corresponding strands of γ continue to two adjacent intersection points in some interval $[p_j, p_{j+1}]$ where $j \geq i$, we already have a tetrahedron in the part of the B piece above them: the convex hull of the edge between them (this edge is unique – both intersections with the branch-line have been identified) and the edge $[x_i, x_j]$. As previously noted, there are too many such tetrahedra to directly yield an adequate volume bound.

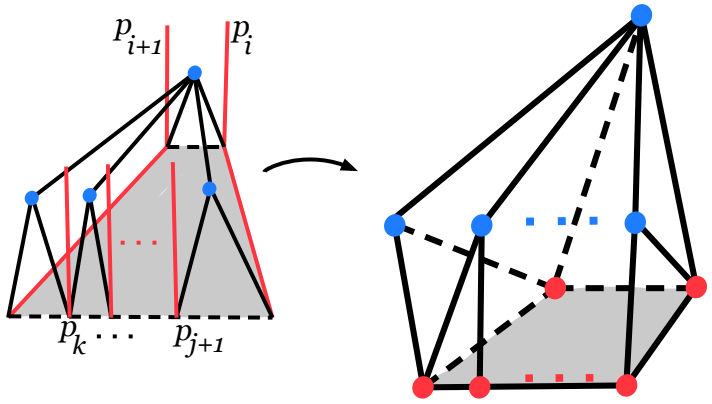

FIGURE 7. The figure on the left shows the polyhedron between two segments of γ that split before returning to the branch-line. The grey portion of the polyhedron is the previously triangulated polygon lying within the template's ear. The figure on the right shows the same polyhedron after taking γ to infinity.

A strand of γ passing through p_i upon entering the B piece from the x ear corresponds to a power of y in the coding of γ. This power determines how many times the strand will wind around the y ear before exiting the B piece. If we denote this power of y by m_{j_i}, this strand will wind around the y ear as a sequence of m_{j_i} segments of γ separated by intersections with the branch-line. We denote these γ segments by $\gamma_1, \ldots, \gamma_{j_i}$ and observe that each of them is adjacent to at most two rectangles within the template. Indeed, a segment is adjacent to a rectangle whenever it does not split from the γ arc crossing the branch-line next to it on the left or the right before returning to the branch-line. In such a situation, starting at p_i and on either side of the γ segments, there may be a chain of consecutive rectangles separated by the branch-line and adjacent to the segments γ_1, γ_2,... γ_{j_k}. This chain terminates just before a splitting of segment γ_{j_k} and its neighbouring γ arc if there is one, or at the intersection with qq'. Thus, on each side, we have a chain of at most m_{j_i} rectangles and above each rectangle we have a tetrahedron (as indicated in the previous paragraph). Recalling that the branch-line edges of these rectangles are all identified as a single edge (see Section 3.1.3), the rectangles in the chain all have this edge

in common. Therefore, by Proposition 3.4, the tetrahedra belonging to this chain contribute at most $C\log(m_{j_i})$ to the volume of $T^1M \setminus \gamma$.

Note that once γ_{j_k} does split from a neighbouring γ arc, then the region above the splitting was already triangulated (see Figure 7). Moreover, the regions in the template that start on the branch-line between the endpoints of the splitting segments were also previously taken care of as they are adjacent to γ segments that start at points $p_j, \ldots, p_k$ that lie between these endpoints. As such, it only remains to triangulate the region above the plane section parallel to the y ear containing the x_i's which are not connected to o_2 by an edge.

Observe that this plane section is divided into a number of triangles and squares that is linear in n_γ. Upon subdividing the squares into triangles, we obtain a planar region subdivided into a number of triangles that is also linear in n_γ. We can then triangulate the upper region by coning these triangles to o_2. To conclude, the total volume of the B piece is thus given by

$$A(n_\gamma + \sum_{i=1}^{n_\gamma} \log(k_i))$$

for an appropriate constant $A > 0$. This bound is at most linear in $|\gamma|$ by Lemma 2.1. □

4. An upper bound for arbitrary hyperbolic surfaces

Having established Theorem 1.1 for the modular surface, we now obtain its general form in three stages. First, we show that the claim is topological, in that it is independent of the choice of hyperbolic structure on a surface:

Lemma 4.1. *Let g and h be two hyperbolic structures on the surface S. If Theorem 1.1 holds for (S, g), then it also holds for (S, h).*

Proof. This follows immediately from the fact that the metrics g and h on S are bi-Lipschitz. Indeed, any Dirichlet fundamental domain is the union of a set of cusp neighbourhoods and a compact complement, so we can take any diffeomorphism which is an isometry near the cusps. Furthermore, we can take the diffeomorphism to be homotopic to the identity. In that case, the diffeomorphism will preserve the isotopy classes of curves. Since every essential closed geodesic is the shortest curve in its isotopy class, it follows that the ratios of lengths of the geodesics corresponding to each such isotopy class in the two hyperbolic structures are uniformly bounded. □

Remark. The supremum of those ratios is exactly the infimum of the Lipschitz constants of maps between the structures; see [19, Thm. 8.5]. Control of the ratios also follows from the bounds in [14].

Next, we show that the claim passes to covers:

Proposition 4.2. *Let $\hat{S} \to S$ be a degree d regular covering of hyperbolic Riemann surfaces. Suppose Theorem 1.1 holds for S with constant A_S. Then it also holds for $\hat{S}$ with the constant $d \cdot A_S$.*

Proof. Given a filling set γ of periodic geodesics on $\hat{S}$, let γ' be the union of all translates of γ under the group of deck transformations and let δ be the projection of γ (and γ') to S. Then, for each component (simple closed geodesic) δ_i in δ, γ must contain one of the lifts of δ_i to $\hat{S}$, which are themselves closed geodesics of length at least $l_S(\delta_i)$. It follows that $l_S(\delta) \leq l_{\hat{S}}(\gamma)$ and, hence, that

$$\begin{aligned} \mathrm{Vol}(T^1\hat{S} \setminus \gamma) &\leq \mathrm{Vol}(T^1\hat{S} \setminus \gamma') = d \cdot \mathrm{Vol}(T^1 S \setminus \delta) \\ &\leq d \cdot A_S \, l_S(\delta) \\ &\leq (d \cdot A_S) \, l_{\hat{S}}(\gamma). \end{aligned}$$

□

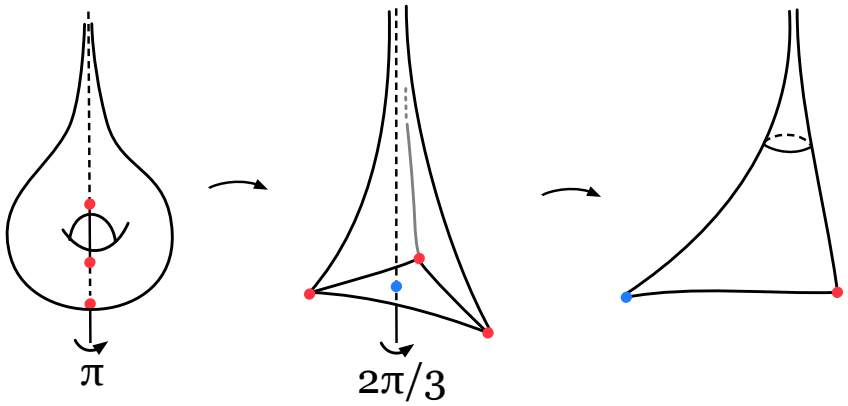

FIGURE 8. The cover of M by a once punctured torus.

Finally, we use ideas of Brooks on adding punctures to hyperbolic surfaces to deduce the general case:

Proof of Theorem 1.1. We have seen that the Theorem holds for any surface which is a cover of the modular surface. This includes the

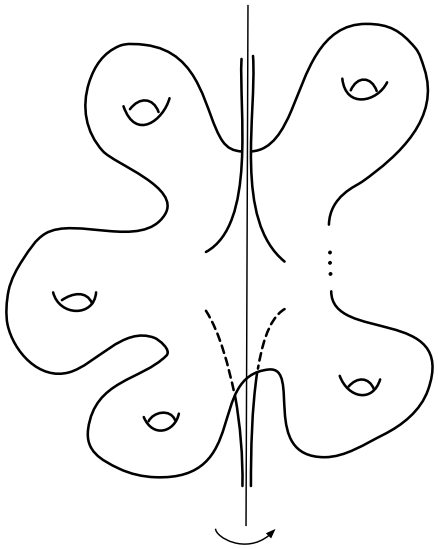

FIGURE 9. The cover of the twice punctured torus by a twice punctured genus g surface.

once-punctured torus (see Figure 8) and its double cover, the twice-punctured torus. (To construct this double cover, arrange the two punctures symmetrically about the axis of revolution of the torus and apply a rotation by π.) In the same manner, there is a $6k$-fold covering map from a k-punctured torus to M. Next, a rotation through the two cusps in a symmetrically arranged twice-punctured surface of genus g (see Figure 9) yields a cover of the twice punctured torus, and thus of M. The same construction gives a covering map from $S_{g,2+kg}$ (the surface of genus g with $2 + kg$ punctures) to M by permuting the punctures and mapping to the $k+2$ punctured torus. Lastly, recall that the thrice punctured sphere is a 6-fold cover of the modular surface. The same strategy now produces a $6(n-3)$-fold covering map from a punctured sphere $S_{0,n}$ with $n > 3$ to M by arranging an axis through two of the punctures and using a rotation acting transitively on the rest of the punctures.

We are now ready to consider a collection of periodic geodesics γ on an arbitrary hyperbolic surface $S_{g,n}$. Since we have already dealt with the case where $n \equiv 2 \bmod g$, we just need a way to increase the number of punctures by k (chosen such that $n + k \equiv 2 \bmod g$) to obtain a hyperbolic surface $S_{g,n+k}$ which is a topological cover of M. For this, mark k arbitrary points on $S_{g,n}$. By [7] (see especially Lemma 3.1), there is a choice of metrics on $S_{g,n}$ and $S_{g,n+k}$ such that

every closed geodesic in $S_{g,n}$ can be isotoped away from the punctures with the resulting closed geodesic in $S_{g,n+k}$ having length bounded by a constant times its length in $S_{g,n}$.

Accordingly, given a family γ of closed geodesics in $S_{g,n}$, let γ' be the isotopic family of curves lying off the cusps. We then have

$$\begin{aligned}\mathrm{Vol}(T^1 S_{g,n} \setminus \gamma) &= \mathrm{Vol}(T^1 S_{g,n} \setminus \gamma') \\ &\leq \mathrm{Vol}(T^1 S_{g,n+k} \setminus \gamma') \\ &\leq C\,|\gamma'| \leq C'\,|\gamma|\ .\end{aligned}$$

Here, the first equality is the topological invariance of the volume and the following inequality is the monotonicity of the volume of hyperbolic three-manifolds under drilling (the circles lying over the punctures). The third claim is the Theorem in the case of $S_{g,n+k}$, and the last step is the linear relation between the length of γ' in $S_{g,n+k}$ and the length of γ in $S_{g,n}$. □

References

1. Colin C. Adams, *Thrice-punctured spheres in hyperbolic* 3*-manifolds*, Trans. Amer. Math. Soc. **287** (1985), no. 2, 645–656. MR 768730 (86k:57008)
2. Dimitri V. Anosov, *Geodesic flows on closed Riemann manifolds with negative curvature.*, Proceedings of the Steklov Institute of Mathematics, No. 90 (1967). Translated from the Russian by S. Feder, American Mathematical Society, Providence, R.I., 1969. MR 0242194 (39 #3527)
3. Emil Artin, *Ein mechanisches system mit quasiergodischen bahnen*, Abh. Math. Sem. Univ. Hamburg **3** (1924), no. 1, 170–175. MR 3069425
4. Joan S. Birman and R. F. Williams, *Knotted periodic orbits in dynamical system. II. Knot holders for fibered knots*, Low-dimensional topology (San Francisco, Calif., 1981), Contemp. Math., vol. 20, Amer. Math. Soc., Providence, RI, 1983, pp. 1–60. MR 718132 (86a:58084)
5. Alex Brandts, Tali Pinsky, and Lior Silberman, *Volumes of hyperbolic three-manifolds associated to modular links*, In preparation.
6. Martin Bridgeman, *Bounds on volume increase under Dehn drilling operations*, Proc. London Math. Soc. (3) **77** (1998), no. 2, 415–436. MR 1635161 (99f:57019)
7. Robert Brooks, *Platonic surfaces*, Comment. Math. Helv. **74** (1999), no. 1, 156–170. MR 1677565 (99k:58185)
8. Pierre Dehornoy, *Geodesic flow, left-handedness and templates*, Algebr. Geom. Topol. **15** (2015), no. 3, 1525–1597. MR 3361144
9. Patrick Foulon and Boris Hasselblatt, *Contact Anosov flows on hyperbolic 3-manifolds*, Geom. Topol. **17** (2013), no. 2, 1225–1252. MR 3070525
10. David Fried, *Transitive Anosov flows and pseudo-Anosov maps*, Topology **22** (1983), no. 3, 299–303. MR 710103 (84j:58095)
11. Étienne Ghys, *Knots and dynamics*, International Congress of Mathematicians. Vol. I, Eur. Math. Soc., Zürich, 2007, pp. 247–277. MR 2334193 (2008k:37001)

12. François Guéritaud, *On canonical triangulations of once-punctured torus bundles and two-bridge link complements*, Geom. Topol. **10** (2006), 1239–1284, With an appendix by David Futer. MR 2255497 (2007g:57009)
13. Dubi Kelmer, *Quadratic irrationals and linking numbers of modular knots*, J. Mod. Dyn. **6** (2012), no. 4, 539–561. MR 3008409
14. Anna Lenzhen, Kasra Rafi, and Jing Tao, *Bounded combinatorics and the Lipschitz metric on Teichmüller space*, Geom. Dedicata **159** (2012), 353–371. MR 2944537
15. John Milnor, *Hyperbolic geometry: the first 150 years*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 1, 9–24. MR 634431 (82m:57005)
16. Peter Sarnak, *Linking numbers of modular knots*, Commun. Math. Anal. **8** (2010), no. 2, 136–144. MR 2587769
17. Caroline Series, *The modular surface and continued fractions*, J. London Math. Soc. (2) **31** (1985), no. 1, 69–80. MR 810563
18. William P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) **19** (1988), no. 2, 417–431. MR 956596 (89k:57023)
19. ———, *Minimal stretch maps between hyperbolic surfaces*, preprint `arXiv:math.GT/9801039`, 1998.

The University of British Columbia, Vancouver BC, Canada
E-mail address: `mbergeron@math.ubc.ca`
URL: `http://www.math.ubc.ca/~mbergeron`

E-mail address: `tali@math.tifr.res.in`
URL: `http://www.math.tifr.res.in/~tali/`

E-mail address: `lior@math.ubc.ca`
URL: `http://www.math.ubc.ca/~lior/`